

ISHM Decision Analysis Tool: Operations Concept

Lilly Spirkovska

NASA Ames Research Center, MS 269-3, Moffett Field, CA 94035-1000

Abstract

The state-of-the-practice Shuttle caution and warning system warns the crew of conditions that may create a hazard to orbiter operations and/or crew. Depending on the severity of the alarm, the crew is alerted with a combination of sirens, tones, annunciator lights, or fault messages. The system uses data such as temperature, pressure, flow rates, and switch positions to determine whether there is an alarm situation. Fault messages directly announce anomalies in the sensed parameters. The combination of anomalies (and hence alarms) indicates the problem. Determining what problem a particular combination represents is not trivial. Even with much training, it is sometimes difficult to determine the root cause of a set of alarms. Integrated System Health Management (ISHM) systems are being developed to address this difficulty. In many situations, an automated diagnosis system can help the crew more easily determine an underlying root cause. Due to limitations of diagnosis systems, however, it is not always possible to explain a set of alarms with a single root cause. Rather, the system generates a set of hypotheses that the crew can select from. The ISHM Decision Analysis Tool (IDAT) assists with this task. It presents the crew relevant information that could help them resolve the ambiguity of multiple root causes and determine a method for mitigating the problem. IDAT follows graphical user interface design guidelines and incorporates a decision analysis system. I describe both of these aspects.

1 Introduction

The goal of Integrated System Health Management (ISHM) is to ensure system functionality. The field encompasses the set of activities that are performed in order to identify, mitigate, and resolve faults with a system[1]. A fault is the root cause of an anomaly, and an anomaly is a detectable undesirable state. Faults may or may not lead to failure. ISHM systems compute and present information in a way that makes it easier for humans to understand the operation of systems and also assist humans in dealing with complexities that may arise. The ISHM Decision Analysis Tool (IDAT) is a prototype of one such ISHM system.

The mission for IDAT is to develop and provide a demonstration of an operations concept for assisting a spacecraft crew in making informed optimal decisions in the

face of uncertainty. IDAT can assist either space-based astronauts or ground-based Mission Control Center (MCC) controllers. Thus, in this paper, “crew” means either the on-board crew or the ground-based crew. Further, IDAT can assist the crew in making a mitigation decision either when there is an uncertain cause for an anomaly or when there is variability in the situational context, as explained below.

I begin by describing the domain for the IDAT prototype and the current state of practice for spacecraft ISHM. I then describe the IDAT tool, first from the user’s viewpoint and then from an implementation viewpoint. I conclude with a partial list of future research that is required in order to improve a crew’s task of ensuring spacecraft system functionality.

2 Domain

The domain for the IDAT demonstration is the Shuttle Reaction Control System (RCS) propulsion system. The RCS is located in three separate modules on the Shuttle, as shown in Figure 1. The forward module is in the nose area, forward of the cockpit windows. The left and right modules are collocated with the Orbital Maneuvering System (OMS) in the left and right OMS pods, near the tail of the vehicle. The 16 forward jets, 14 left jets, and 14 right jets control the motion of the Shuttle. Each jet is permanently fixed to fire in a particular direction: up, down, left, right, forward, or aft. The selective firing of individual or combinations of jets provides rotational (about an axis) or transitional (along an axis) movement. Rotational movement is used for attitude control, whereas translational movement is used primarily for velocity changes.

The RCS can be used alone to provide attitude control, or in combination with the OMS or Main Propulsion System (MPS). The RCS is used throughout a typical Shuttle mission. During the ascent phase (known as Operational Sequence (OPS) 1 in the flight software), the RCS is used for external tank (ET) separation in nominal situations or to assist the MPS with additional roll control during off-nominal (single-engine) situations. During the on-orbit phase (OPS 2), the RCS provides attitude control for basic orbiting as well as during rendezvous and proximity operations. Finally, during the de-orbit phase (OPS 3), the RCS is used to orient the Shuttle prior to de-orbit burn to optimally reenter Earth’s atmosphere. (Figure 2 shows the operational sequences of a Shuttle flight.) A minimum number of jets are required to be operational for each of these activities. If this minimum is not achieved, for example, the ET and orbiter could collide during tank separation.

The RCS jets use monomethyl hydrazine as the fuel and nitrogen tetroxide as the oxidizer. These two propellants are hypergolic, which means that ignition occurs spontaneously upon fuel-oxidizer contact, thus eliminating the need for an ignition device. Although this approach increases system reliability, the volatile, corrosive,




Figure 1. Location of Reaction Control System (RCS) on Shuttle




Figure 2. Operational sequences (OPS) of Shuttle flight software and the major mission modes (MM) within a sequence

and poisonous nature of these propellants adds an operational risk if any leakage occurs.

Each RCS module contains a collection of jets, a fuel tank, an oxidizer tank, and two helium (He) tanks, along with associated feedlines, manifolds, and other supporting equipment. Propellant flow (fuel and oxidizer) to the jets is normally maintained by pressurizing the propellant tanks with helium. Figure 3 shows a schematic of the fuel side of an RCS module. The same equipment is duplicated for the oxidizer side, as well as for the fuel and oxidizer sides of the other RCS module.

Because the aft RCS modules are collocated with OMS engines and both use the same types of propellant, it is possible to interconnect (*i'cnct*) the two systems to allow for propellant sharing. It is also possible to share propellant by establishing a crossfeed (*xfeed*) between the two aft RCS modules.




Figure 3. Schematic of the fuel side of an RCS module. Similar components exist for the oxidizer side.

3 Background

In the current state of practice, the crew must monitor the RCS (as well as the other Shuttle systems) by watching and scrutinizing spreadsheets of numbers. The state-of-practice display for the summary of the status of the OMS and RCS systems is shown in Figure 4. The top left and right upper half of the figure shows the quantities, pressures, valve positions, and fuel injector temperatures of the OMS, whereas the area below the topmost horizontal lines shows the propellant, manifold, and jet status of the RCS. The crew also has access to an RCS Redundancy Management (RM) status display, shown in Figure 5. The RCS RM is a series of GPC (on-board general-purpose computer) processes that detect, identify, and monitor such items as failed jets (failed on, failed off, or failed leaking) and manifold valve status. The results of the RM processes are displayed on the aforementioned display. The figure shows the data for the forward RCS pod (signified by the “*” next to “RCS FWD 1”). As one can speculate, it requires considerable training to learn how to interpret this display. Massive amounts of data about the system are presented, but it is not trivial to extract the desired information.

1031/ /019				GNC SYS SUMM 2				5 000/00:03:34			
								BFS 000/00:00:00			
OMS AFT QTY		L		R		OMS		L		R	
		OXID 30.1		30.1		TK P		HE 3700		3730	
		FU 30.1		30.1				OXID 282		280	
		FU INJ T		79 79				FU 280		277	
								N2 TK P		2190 2180	
								REG P		326 329	
								P VLV OP		OP OP	
								ENG IN P			
								OXID 280		281	
								FU 277		280	
								VLV 1 - 2		- 2	
								2 0		0	

ponents, such as valves, tanks, and flow lines, around a spatial arrangement that emulates key aspects of the underlying system architecture. In addition, the designers incorporated luminance and color-coding schemes into the graphics so as to provide “at-a-glance” information concerning the current system configuration and operating mode [2,3].

Improved displays make a significant difference in helping the crew extract relevant information about a system. Nevertheless, in general, people tasked with monitoring a system tend to get tired and decrease their vigilance, especially over a long duration. They also tend to become complacent in situations where usually nothing breaks. Computers are better suited to such repetitive work.

With this in mind, the underlying concept of operations for IDAT is that automation is used where appropriate for the monitoring task. Additionally, other ISHM technologies are integrated with automated monitoring to decrease the crew’s workload and to assist them in dealing with system complexity. In the case of IDAT, we specifically refer to complexity due to uncertainty and variability.

Uncertainty arises when a set of anomalies can be explained by multiple root causes and it is not possible to distinguish the true root cause. This can occur for a number of reasons. First, the leading techniques for automated diagnosis rely on building a model of the system and comparing the predictions from the model to the observations of the real system. Due to the complexity of spacecraft systems, it is challenging to build high-fidelity models. Second, the complex interactions between components make it difficult to apply a variety of other diagnosis techniques, due to the limited ability of the system experts to think of all possible problems in advance and to determine how each problem could affect other components or systems. Third, it is impractical to place a sufficient number of sensors throughout the system to be able to disambiguate between some root causes, and the sensors themselves tend to be less reliable than the systems being observed. In these cases of ambiguous root cause, IDAT works in tandem with the crew to help them select a strategy for mitigating the fault. It follows the task-oriented approach by presenting germane information that enables the crew to make a decision by coordinating their knowledge or constraints with that provided by the automation.

IDAT also assists the crew in dealing with complexity due to variability, by which we mean that multiple mitigation procedures are possible for a given root cause. In these cases, the procedure to follow is determined by the situational context. For example, the crew’s best option may be to select the Isolate procedure to further refine the cause of a left RCS leak when the leak is small, there is adequate troubleshooting time available, and the propellant reserves are plentiful. In contrast, if the leak is large, the right RCS is feeding off the left because of a previous leak, the propellant level is approaching critical levels, and only five minutes have elapsed since launch, the best option may be to select the Abort procedure. Many of the standard on-board procedures begin with a list of constraints that must be fulfilled to use that

procedure and pointers to other procedures for cases where those constraints are not met. IDAT assists by automating the evaluation of constraint satisfaction as feasible given the available sensors. It provides the crew with a list of constraint-satisfying options, as well as the advantages and drawbacks of each option.

The following section describes the IDAT tool. I introduce the elements displayed on it sequentially and reference the labeled areas of Figure 6.

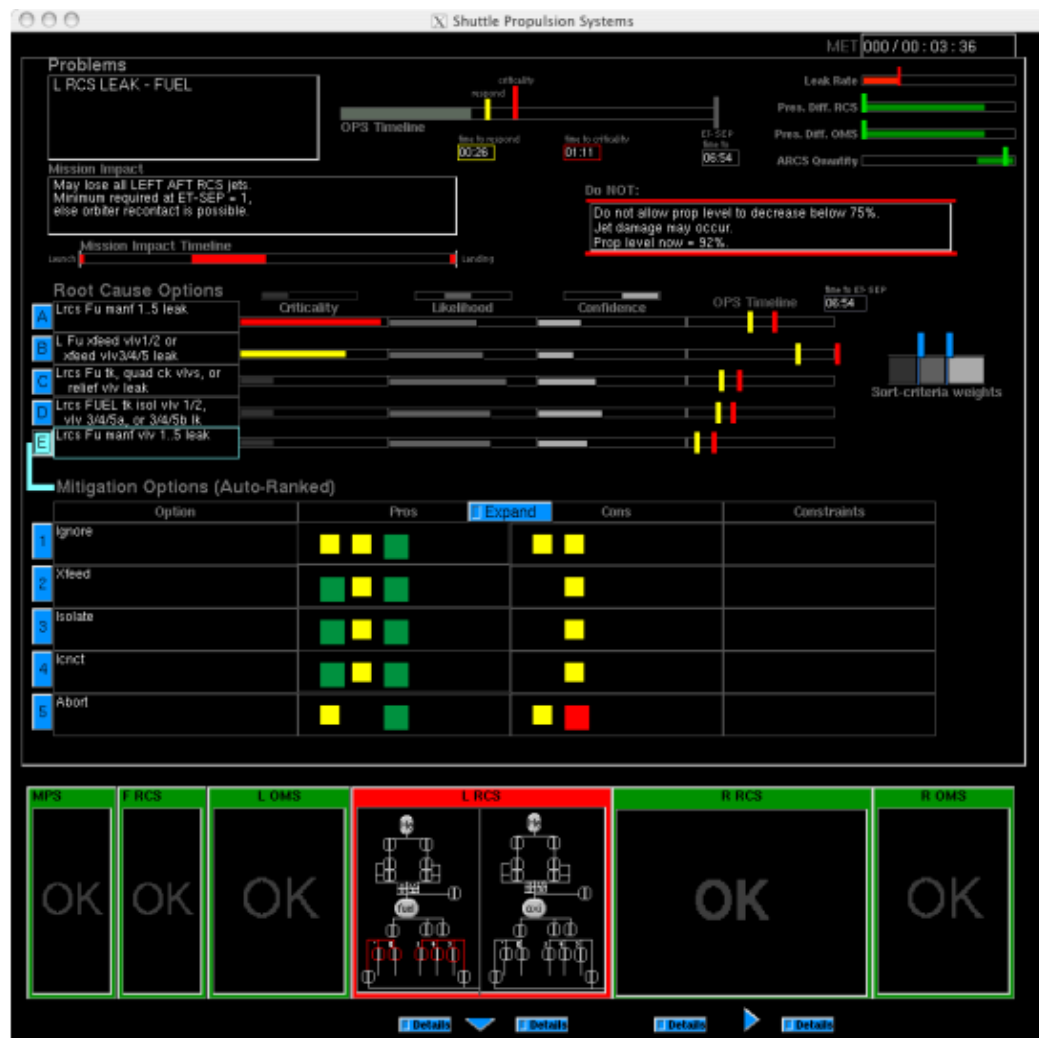


Figure 6. IDAT main display

4 IDAT

The operational concept of IDAT is that the system health state determination is automated as much as possible and the crew are alerted only when a problem they need to assist with is detected.

When the propulsion systems' monitored parameters are nominal, the IDAT display shows only the mission elapsed time (MET, shown in the upper right of Figure 6) and a status summary for each system. The six systems shown are the main propulsion system (MPS), the forward RCS (F RCS), the left OMS (L OMS), the left and right RCS (L RCS and R RCS), and the right OMS (R OMS). They are arranged as shown because the aft RCS and OMS share a pod on each side of the shuttle, the F RCS is independent, and the MPS is not of concern once in orbit and can be deleted from the display without changing the order of the other systems. The size of each status summary area can be adjusted to reflect its importance during a particular phase of flight. The enlarged RCS summary areas reflect the choice of domain for the prototype system. When all monitored parameters are nominal, the status is shown as "OK" and the border is color-coded green.¹

When a problem is detected, the border changes to red and a summary of the problem (e.g., "LEAK - FUEL") appears. (Figure 9 shows this view whereas Figure 6 shows the expanded version of the summary area.)

Simultaneously, the remaining IDAT display appears. The upper-left section, labeled "Problems," shows the current list of problems using phrasing that mimics the phrasing used in the current Caution and Warning system. In our example, the problem is "L RCS LEAK - FUEL," meaning a leak has been detected in the left RCS fuel propellant leg.

An automated diagnosis system coordinates with the monitoring system to determine the root cause of the problem. As discussed previously, determining the root cause is not always easy. Humans frequently create systems whose behavior, particularly in fault cases, is so complex that even the creators cannot fully predict it. The root cause of a fault may not be any single component, but rather a possibly complex interaction between components. Further, because of a variety of factors, such as low-fidelity models or inadequate number of sensors, the diagnosis system is not always able to determine a single root cause, but rather generates a set of possible root causes that explain the anomalous behavior. Such is the case for this problem. There is an inadequate amount of data to discern the location of the leak. IDAT thus displays the top-ranked root cause options (in the middle of the display, labeled "Root Cause Options" and preceded by the letters "A" to "E").

The root causes have characteristics associated with them that provide additional information to the crew. We selected three characteristics that we believe are useful. Others may also be important and closer interaction with the eventual users is an important step in defining a complete set of characteristics. The current set of characteristics includes criticality, likelihood, and confidence, defined as follows:

¹ To increase the crew's comfort level that the automation is still working, rather than displaying just an "OK," it may be desirable to introduce a heartbeat-like display that shows the current distance from perfection as, for example, an intensity level of the "OK" or as symbols that grow/shrink to represent the distance.

- Criticality refers to a scale of possible ramifications of a root cause (a.k.a. fault). The highest criticality is assigned to faults that could result in loss of human life, while the lowest criticality is assigned to nuisance faults that lead to only a slight degradation of current or potential future performance. The criticality of a fault is due both to the misbehaving components and to the current phase of flight.
- Likelihood refers to the prior probability that the components (not sensors) called out in the root cause may fail. The prior probability considers such factors as the mean time between failure (MTBF) metric, as well as any previous anomalous behavior of a component during the current mission.
- Confidence refers to the level to which the diagnosis system believes the behavior is explained by the implicated components. For example, model-based diagnosis compares observations of the real system with the predictions from a model. In the case of a fault, discrepancies between the observed behavior and the predicted normal behavior occur. These discrepancies are then used to identify (diagnose) the fault. In practice, the observations from the system (relayed by sensors strategically placed throughout the system) may not be received in time to reach a hypothesis. These missing alarms caused by a time-out event that occurs before a sensor reading is received decrease the confidence of the diagnosis system in its stated hypothesis. It is also conceivable that the sensor has failed and is providing an incorrect reading. In other words, the confidence characteristic refers to what the sensors are currently (and recently) saying, whereas the likelihood refers to the history of the system components. Note that if sensors and components are treated equally, it may be possible or desirable to combine the likelihood and confidence values into a single value.

The three rectangles to the right of the text of each root cause encode the values for the three characteristics. The value of each characteristic is shown by the fill amount within its associated rectangle. Three distinct hues of gray represent the three characteristics: the darkest gray is associated with criticality, medium gray is associated with likelihood, and the lightest gray is associated with confidence. To draw extra attention to higher criticality values, the fill for criticality uses gray to represent lower criticality values, but adds color to represent higher criticality values, using yellow to represent moderate criticality, and red to represent high criticality.

A weighted summation of the values of the three characteristics determines the ranking of the root cause options and thereby the order in which they are displayed. The crew can adjust the weight afforded each characteristic to better reflect personal ranking of the importance of each characteristic. IDAT will automatically resort the list of root causes as the crew adjusts (moves) the blue sliders of the “Sort-criteria weights” area. The three subrectangles defined by the two blue sliders map to the three characteristics. We use several cues to help the crew associate a subrectangle with an individual characteristic. First, the subrectangles use the same order as used in the “Root Cause Options” area: criticality on the left, likelihood in the middle, and confidence on the right. Second, the same color represents an individual

characteristic in both areas. Third, the sliders change the size of the subrectangles in the “Sort-criteria weights” area and within the rectangles above each characteristic label in the “Root Cause Options” area.

In addition to the value of these characteristics, we believe it is important for the crew to know the amount of time until a mitigation option must be started and the amount of time until a particular root cause leads to a critical situation (i.e., failure). These two times are associated with each root cause option and shown next to the characteristics as “OPS Timeline.” The time to criticality (“ttc,” shown by the red bar) depends on the mission mode (i.e., the functions of the system during different phases of flight) and encodes how long it takes until critical ramifications result from the root cause. The time to respond (“ttr,” shown by the yellow bar) takes into consideration the amount of time required to perform a mitigation activity.

Any one of the root causes presented by IDAT could be the true root cause. The automated system does not have enough information at the current time to distinguish between them, so it is up to the crew to do so. To assist with this process, we recall some of the parameters from the original state-of-practice monitoring system, namely the parameters or cues that are relevant for this particular problem. As shown in the upper-right corner of Figure 6, rather than the numeric display format, we present the relevant cues graphically. Also, to ease detection of out-of-range conditions, we use color coding to supplement a parameter’s value with information about its relevant position with respect to a hard limit threshold (redline) and acceptable range that is specific to each parameter. If the parameter’s current value is within the acceptable range, the bar is color-coded green. When it approaches the limit threshold, the bar and surrounding rectangle both turn yellow. When the value reaches the limit threshold, the bar and surrounding rectangle both turn red. Note that the bar continues to update its position to reflect updates in sensed values.

Beside the relevant parameters, we also provide the crew with timing information. This is shown within the “OPS Timeline” rectangle between the relevant cues and “Problems” area at the top of Figure 6. The OPS timeline shows the mission elapsed time within that operational sequence, the time available to respond, time to criticality, and time until the end of the current operational sequence. The elapsed time is shown by gray filling of the rectangle. The timeline also includes a countdown timer until the end of the sequence, and the activity that signifies that point (“ET-SEP,” or external tank separation, in our case). The time to respond refers to the latest time a response can be started to mitigate the highest-ranked root cause option shown in the “Root Cause Options” section, i.e., the root cause with label “A,” and is shown graphically by the yellow “respond” line and “time-to-respond” text value. The time to criticality refers to the time until the situation becomes critical if the highest-ranked root cause (i.e., root cause “A”) is the true cause of the problem, and is shown graphically by the red “criticality” line and “time-to-criticality” text value.

To help the crew assess the criticality, IDAT displays both text and a graphical presentation of the impact this problem has to the mission (from the start of the current operational sequence to landing). This information is shown in the “Mission Impact” area located below the “Problems” area.

IDAT displays any critical actions that the crew absolutely must *not* take for that type of problem, as shown in the “Do NOT” area.

To visualize the root causes, the crew can select the “expand” option (shown as a sideways triangle similar to its use in many mainstream computer applications) to replace the text with a summary representation of the schematic for the affected system, as shown within the “L RCS” status summary area. The components of each root cause are highlighted in red to assist with visualization of the part of the system that is being implicated within the text of the selected root cause. The crew can select any of the other root causes to visualize the other options. Note that throughout the display, anything shown in blue is selectable. The crew can also bring up a more detailed version of the schematic that also shows numeric values for associated parameters. The display, obtained by selecting the “Details” button, is shown in Figure 7. A root cause may implicate components from multiple systems. That is, a root cause may state that the symptoms can be explained by a combination of a fault of a component in the left RCS and a component in the right RCS. To visualize this situation, the crew can “expand” the status summary area for both the left and right RCS. It is also possible to view the “Details” windows for both modules side by side.

At this point, IDAT has presented all the relevant information (barring additional information that may be recommended from domain experts during follow-on work) that could help the crew resolve the ambiguity of multiple root causes and begin to determine a method for mitigating the problem. The most important action for the crew is to ensure system functionality. This is aided by the proper diagnosis of the root cause, but can be accomplished even when the root cause cannot be determined.

To assist the crew with the task of determining a mitigation strategy, IDAT incorporates a decision analysis system. We use a COTS (commercial off-the-shelf) decision network tool, Netica, to automatically rank available mitigation options based on the inputs of relevant cues and knowledge about the effect of a mitigation option on resources. The ranked set is displayed in the “Mitigation Options” area, along with the effect on resources (“Pros” and “Cons”) and any associated flight rules (“Constraints”). The pros and cons can be shown either as color-coded rectangles (as shown in Figure 6) or text (as shown in Figure 8). The effects on resources are classified as benefits (pros) and drawbacks (cons) of vehicle damage, crew resources, and propellant resources. These three factors were used in the prototype implementation, but can be modified to incorporate domain expert guidance.

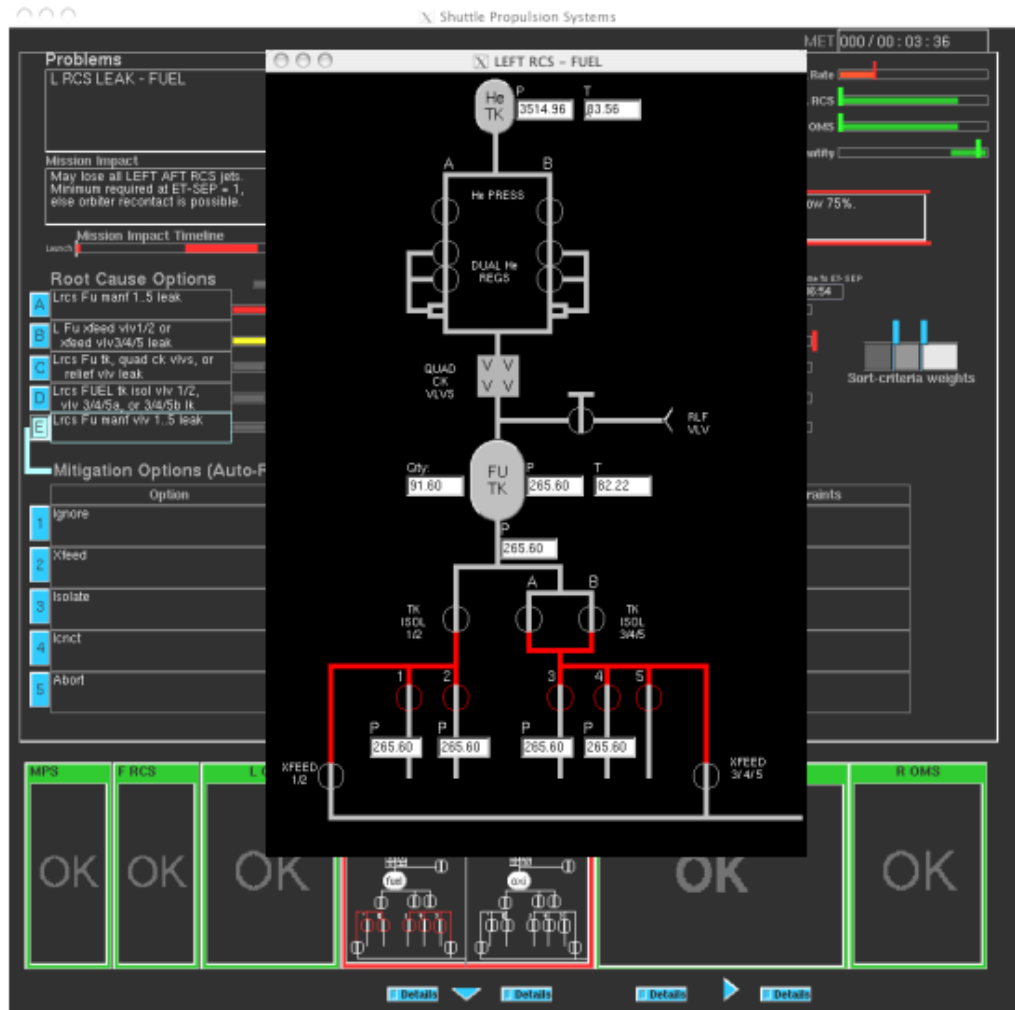


Figure 7. IDAT detailed schematic display

The numbers (1 to 5) to the left of each mitigation option enable automatic access to relevant procedures that the crew needs to accomplish. The crew may also choose to review details of an option to help them make a better selection. We have not yet completely implemented this capability, but rather just provided hooks for the process. In particular, clicking on one of the numbered blue boxes currently brings up only the titles for the relevant procedures, as shown in Figure 9. Potentially, it may be more useful to display both the title and the complete procedure. In Section 7, I discuss introducing automation into procedure execution to further reduce the crew's workload.

In summary, we created an operational concept in the embodiment of a tool (IDAT) that works in tandem with the crew. It is a decision support tool that assists the crew in determining an appropriate mitigation strategy when confronted with a problem with an uncertain cause. IDAT provides much germane information in the form of context and relevant cues to augment the crew's knowledge and help them select the best action. In addition to helping the crew deal with uncertainty, it can also




Figure 8. IDAT expanded pros and cons within the *Mitigation Options* area

be valuable in determining a mitigation strategy in situations with variable context. It integrates a variety of automated systems (monitoring, diagnosis, and decision support) with the engineering systems of a complex vehicle like the Shuttle or CEV (Crew Exploration Vehicle, the replacement for the Shuttle slated to fly around 2012), but does it in a manner that keeps the human in the center of the process.

In the following sections, I provide more detailed descriptions about the underlying pieces of IDAT. I begin by describing the scenario used to demonstrate the concepts, the architecture of the tool, and the details for each subsystem. I conclude with a discussion of future work necessary to further develop the operational concept.




Figure 9. IDAT expanded pros and cons within the *Mitigation Options* area

5 Scenario

The IDAT prototype uses the Shuttle RCS as a demonstration platform. As a illustration case, we introduce a leak in one of the components of the fuel propellant leg of the left RCS. The crew's task is then to determine the location of the leak and decide how to proceed with the mission. A diagnosis system produces from three to five possible root cause options, some with different values for fault criticality and time to criticality. The available mitigation options are to abort the mission, to ignore the leak, to attempt to isolate the leak, to initiate a crossfeed from the right RCS, or to initiate a crossfeed (known as an interconnect) from the left OMS. We can vary the time the leak begins and the size of the leak.

6 Prototype Details

The architecture of the prototype is shown in Figure 10. The crewmember – and thus the graphical user interface (GUI) – is at the center. Connected to the GUI are the data system, the monitoring system, the diagnosis system, and the decision analysis system. All five systems are implemented in C and C++ and run under Linux. The data system runs as an independent process, writing its output to a file. The other four systems are combined into a process and receive input from that file as well as from the user/crewmember.




Figure 10. IDAT architecture

The purpose of the GUI is to integrate the output from the other systems, call the appropriate systems under the appropriate conditions, display relevant information to the crew, and enable the crew to interact with that information. The graphics portion is implemented in OpenGL.

The purpose of the data system is to generate a data stream for concept demonstration. It uses the RCS parameters from the STS-114 data (Shuttle Discovery mission, July, 2005). Because there wasn't an actual leak during the flight, we had to simulate one. The biggest challenges were establishing how to simulate a leak in each of four possible configurations (normal, secured, crossfeeding, and partial crossfeeding) and determining how to transition between the eight configurations (the four above crossed with ok or leaking). It was also tedious to map the desired parameter to the parameter identification (MSID) used by the Shuttle program. The data system appears on the user interface in the relevant cues area, the OPS timeline area, and the actual state of schematic components. It should also provide data for the "Mission Impact Timeline" but currently does not. Instead, we use representative data for that portion.

The purpose of the monitoring system is to monitor incoming data for anomalous behavior. For the IDAT prototype, we used a substitute system that only looks for a leak. When it detects a leak, it informs the GUI, which causes the main area to appear. The monitoring system's output appears in the "Problem" area.

The purpose of the diagnosis system – DSS for diagnosis system substitute – is to identify root causes of an anomalous situation that the monitoring system detected. This too is a substitute system. It uses a model-based approach. The model is a mental model, the computation uses a biological neural network, and the computation results are entirely hand coded. DSS provides root causes in the various configurations where the leak could arise and considers the state of both RCS modules in its determination of root cause. Further, DSS sets the likelihood and confidence values to semi-realistic values. However, for full demonstration of the GUI, we adjusted some of the values to better show the color-coding and sorting features. DSS also sets the criticality, time-to-respond, and time-to-criticality values. In a full IDAT implementation, these values are expected to originate in a different type of system that considers not just the root cause but also any redundancies available, the

mission profile, and other relevant factors. The DSS output appears on the GUI in the “Root Cause Options” section and as the state of schematic components per selected root cause (both in the status summary area and in the detailed schematics).

Finally, the purpose of the decision analysis system is to rank mitigation options for highest utility (goodness). IDAT uses Netica to compute utility for each of the five mitigation options based on a basic model of five inputs (entered into *chance nodes* or *nature nodes* because they model the nature or reality of the world – the likelihood of it being in its possible states [4], also known as *current-state nodes*) and three expert-knowledge nodes (also chance nodes but referred to as *outcome-state nodes*).²

The Netica model created for IDAT is shown in Figure 11. The current-state nodes encode the situational context via the parameters leak rate, time to criticality, pressure difference between the two aft RCS modules, propellant quantity, and state of the other RCS module. Each of these is specific to the leg with the propellant leak. The single *decision node* of the network, named *Options*, displays the calculated utility for each of the five possible mitigation options when a leak is detected. These options are to ignore the problem, attempt to isolate it, establish a crossfeed to use propellant from the other RCS, interconnect to the OMS to use its propellant, or abort the mission. Each of these options comes with issues that must be considered. These issues – resources used, vehicle damage that may occur, and crew factor, including such things as the workload of the crew – are coded in the values assigned to the expert-knowledge or outcome-state nodes of the network. Finally, the *utility node*, simply named *Utility* in the model, computes the answer to the question “Given the current states, how happy will I be with the decision, considering that it will likely result in the expected future states?”

The biggest challenge with the decision analysis system is parameter optimization. The conditional probability tables (CPTs) have to be manually initialized to reflect the expert knowledge so the network can better answer the question “What should we do, what is the best choice, under the conditions encoded by the input nodes?” This is a difficult process that requires much adjustment. Also, the model is fairly rudimentary in that the decision network does not incorporate the possible root cause in its view of the situational context. That is, the utility computation will result in the same vector of values regardless of the root cause of the problem. The output of the decision network appears on the UI in the “Mitigation Options” area. The five options are ranked by utility, given the inputs represented directly by or derived from the relevant cues (upper right corner of the display), whereas the pros and cons encode the values computed for the expert-knowledge nodes.

² Unlike current-state nodes, the values for outcome-state nodes are not input as evidence, but rather are computed by Netica to capture the *goodness* of the outcome state for the current-state nodes and the decision that is selected.




Figure 11. Netica model created for IDAT

7 Future Work

Through this small demonstration system, we learned that many issues must be resolved to ensure that IDAT provides the crew with appropriate and desirable assistance in assuring the health of the complex systems required to successfully and optimally³ accomplish their mission. In this final section, I present six categories of issues that need to be resolved to flesh out the operational concept: user interaction, data system, diagnosis system, decision support system, mitigation procedures, and human factors (display) issues.

User Interaction

- Involve the eventual users in the development of the operations concept to incorporate the domain expertise necessary to evaluate each proposed concept.
- Resolve under what conditions IDAT appears.
- Consider benefits of having IDAT appear immediately when a problem is detected versus delaying its appearance based on the crew's other work. Investigate how to determine and track crew's current tasks and relative priorities.
- Investigate benefits of hiding the problem if a automated mission manager determines the problem will not have an effect on this mission. (Requires trend information from a prognostics system.)

³ There is a saying in aviation: "A good landing is one you can walk away from. A great landing is one that leaves the plane usable for another flight." Similarly, a successful mission is one that does not result in loss of human life. An optimal mission is one that also accomplishes all the major goals and leaves the vehicle in a state that requires only minimal refurbishment for another flight.

- Resolve if and how to enable the crew to defer addressing a problem.
- Resolve how the crew should interact with the procedures that appear when a mitigation option is selected. Possible options include: (1) The computer executes all actions automatically with crew intervention only when automation is not possible. (2) The crew performs each action manually to remain proficient on procedures and locations of switches. (3) The crew may choose some combination of automation and manual action.
- Resolve whether mitigation options should be filtered by root cause, redundancies, mission goals, procedures available, crew's other tasks, etc.
- Determine whether to consider available mitigation procedures in pruning the set of root causes presented to the crew.
- Resolve what happens when additional problems are detected while the crew is still working on the previous problem. Determine what type of new problem (and priority of problem) should cause IDAT to interrupt the crew's interaction with the current problem.
- Resolve what happens when there have been multiple problems associated with a single system. Consider displaying just the most recent problem versus a historical list of problems and their resolution.
- Determine desired interaction for the mission impact timeline.
- Investigate use of other modalities such as aural (including speech) and tactile.

Data System

- Involve domain experts to select a plausible and convincing scenario.
- Incorporate flight rules and other constraints into displayed information.
- Incorporate mission impact information.

Diagnosis System

- Involve ISHM domain experts for information on diagnosis systems current capabilities. Relay information to diagnosis system experts about desired future capabilities.
- Replace the hand-coded diagnosis system with a suitable automated diagnosis system.
- Determine how to compute criticality. Likely need to consider root cause, available redundancies, and the mission.
- Determine how to compute likelihood.
- Determine how to compute confidence, its meaning, the types of diagnosis systems it applies to, and under what situations it can be combined with likelihood.
- Determine whether there are characteristics of a root cause that are important in addition to criticality, likelihood, and confidence.
- Determine how to summarize multiple singletons into a single root cause item.
- Determine whether to merge two root cause options that have a duplicate set of recommended mitigation options.
- Determine whether and how to incorporate a troubleshooting system into IDAT.

The troubleshooting system would enable the crew to interact directly with the diagnosis system to find out the cause of and how to decrease the ambiguity.

- Determine how to incorporate an explanation facility.

Decision Support System

- Expand model to include the possible root cause in its decision computations.
- De-cluster utility values. The current network produces values for utility of the best option and worst option that may vary by as little as a few percent. It may be more useful to a crewmember to receive more polarized advice on the goodness of the possible mitigation options.
- De-cluster pro/con values. Adjust the computation of the values for the expert-knowledge nodes to increase the range of possible values, when appropriate.

Mitigation Procedures

- Incorporate ideas from PRS (Procedural Reasoning Systems, SRI, 1988-1990) for extracting appropriate procedure for a given situation (context-aware procedure retrieval).
- Determine if and how to create procedures on the fly for novel situations.
- Determine if and how to interlink existing procedures to accommodate novel situations.

Human Factors (Display)

- Determine how many root cause options to display based on how many root cause options a person can interpret and extract meaningful information from.
 - If fewer than all options are presented, determine the filtering criteria.
 - If all options are available, but only some are displayed, determine how best to interact with all options.
- Determine how to deal with root cause updates received while the crew is investigating a previously received root cause, say RC1. Consider the case when RC1 is not in the new list and the case when RC1 moves up or down in the new list. Determine how the dynamics of displaying the new list of root causes will affect the crew's analysis or other tasks.
- Establish how cross-subsystem (in general, across multiple subsystems) multiple faults are displayed.
- Determine how to display interactions or interdependencies among root causes.
- Resolve how to visualize interconnect and crossfeed operations.
- Determine how to display double micro-switch valves. RCS uses double micro-switch valves, where each valve produces two sensor values: one for open and one for closed. If a valve is in transition, both values may indicate open (because it has not yet closed).
- Relevant cues are correlated (e.g., leak rate, time to criticality, and quantity). Determine whether to combine correlated cues into a single derived parameter or display each cue separately.

- Time to criticality and time to respond are uncertain estimates. Determine how to relay the probability distribution associated with them to the user.

This is just a short list of issues that should be resolved. Many additional issues will likely arise as IDAT is connected to proper automated systems that provide data, monitoring, and diagnosis services, and as experts are linked into the project.

Acknowledgements

The IDAT project was a collaborative effort of the author with NASA Ames Research Center colleagues David Iverson, Rob McCann, Charles Mott (University of North Dakota), Richard Papasin, Scott Poll, Peter Robinson, and Corinne Ruokangas (RSC). The author also thanks colleagues Miwa Hayashi for her human factors recommendations, Charles Lee for his help in downloading Shuttle data, and Serdar Uckun and Ann Patterson-Hine for the project inspiration and their encouragement.

References

- [1] Stephen B. Johnson. Introduction to system health engineering and management in aerospace. In *First International Forum on Integrated System Health Engineering and Management in Aerospace*, Napa, CA, November 2005.
- [2] J. McCandless, B. Hilty, and R. S. McCann. Upgrades to the caution and warning system of the space shuttle. In *Human Factors and Ergonomics Society 47th Annual Meeting CP*, pages 16–20, Denver, CO, 2003.
- [3] J. McCandless, B. Hilty, and R. S. McCann. Development of new displays for the space shuttle cockpit. *Ergonomics in Design*, 2005. (In Press).
- [4] Stuart J. Russell and Peter Norvig. *Artificial Intelligence: A Modern Approach*. Prentice Hall, Inc., 1995.